



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ :

G06F 12/02

A1

(11) International Publication Number:

WO 00/60470

(43) International Publication Date:

12 October 2000 (12.10.00)

(21) International Application Number: PCT/EP00/02094

(22) International Filing Date: 9 March 2000 (09.03.00)

(30) Priority Data:

9907283.7

31 March 1999 (31.03.99)

GB

(71) Applicant: KONINKLIJKE PHILIPS ELECTRONICS N.V.
[NL/NL]; Groenewoudseweg 1, NL-5621 BA Eindhoven
(NL).(72) Inventor: MORRIS, Steven; Prof. Holstlaan 6, NL-5656 AA
Eindhoven (NL).(74) Agent: WHITE, Andrew, G.; Internationaal Octrooibureau
B.V., Prof. Holstlaan 6, NL-5656 AA Eindhoven (NL).(81) Designated States: BR, IL, IN, JP, KR, European patent (AT,
BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU,
MC, NL, PT, SE).


Published

With international search report.

(54) Title: MEMORY RECLAMATION METHOD

(57) Abstract

A method and apparatus for reclaiming memory space allocated to data structures comprising data objects linked by identifying pointers. A repetitious cycle is performed including the following steps of traversing the data structures to identify and mark active data objects (300), deleting data objects that were marked as garbage during the previous cycle (320), marking data objects not marked as active as garbage (330), and executing the finaliser of any data object marked as garbage having a finaliser (340).



BEST AVAILABLE COPY

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

MEMORY RECLAMATION METHOD

The present invention relates to a method and apparatus for handling
5 stored data objects and particularly, but not exclusively, to the handling of
finalisation for objects in memory compaction and garbage collection
procedures executing in real time in real or virtual memory space of a data
processing apparatus.

10 Garbage collection is the automated reclamation of system memory
space after its last use by a programme. A number of examples of garbage
collecting techniques are discussed in "Garbage Collection: Algorithms for
Automatic Dynamic Memory Management" by R. Jones et al, pub. John Wiley
15 & Sons 1996, ISBN 0-471-94148-4, at pages 1 to 18. Whilst the storage
requirements of many computer programs are simple and predictable, with
memory allocation and recovery being handled by the programmer or a
compiler, there is a trend toward functional languages having more complex
patterns of execution such that the lifetimes of particular data structures can no
20 longer be determined prior to run-time and hence automated reclamation of
this storage, as the program runs, is essential.

Finalisation is a concept used in Sun Microsystems' Java 7 and other
current garbage-collected languages (with Java being a trade mark of Sun
Microsystems) and programming environments, such as Modula-3 and Cedar.
25 Stored data objects created by a program process may no longer be needed
at a point during the execution of the program process. Unneeded data
objects are detected by the environment's garbage collector and marked as
garbage so that the memory occupied by the data object can be reclaimed. A
data object may have an associated finaliser which is to be executed after the
30 object is detected as garbage but before the memory occupied by the data
object is reclaimed. The purpose of this feature is to allow an object to clean
up any other system resources the object has claimed before it is destroyed.

For example, the finaliser for a Java File object would close all the system file handles held by the object.

However, as a finaliser is just a special type of program procedure associated with a data object, with all the power of other program procedures, the finaliser procedure can access and manipulate all data objects accessible from the object being finalised. Therefore, all objects accessible by a finaliser, such as descendant objects accessible from referencing pointers held by the data object, must be explicitly excluded from garbage collection. Furthermore, it is possible for the finaliser method to resurrect any such objects accessible by a finaliser, including the object being finalised itself, by making the object accessible by the program process again. Consequently, a garbage collection procedure cannot delete any objects that are accessible by a finalisable object until its finaliser has executed and the accessibility of the objects has been re-evaluated. In Java and other languages, the possibility of an object repeatedly resurrecting itself is typically removed by stating that the finaliser for each instance is executed only once. This control on finalisation will be assumed herein.

In PC's or workstations, the extra processing and memory load to support finalisation is not usually a problem due to the amount of memory typically available in a PC, although the support will, of course, affect the overall efficiency. In low-memory environments such as set-top boxes, however, support for finalisers can cause problems and even a concurrent or incremental garbage collector may have to halt the program until it has executed some or all of the outstanding finalisers and reclaimed any memory used by them.

In a conventional mark-sweep garbage collector a sweep phase is only executed after a mark phase has finished. Therefore only unmarked objects may be safely treated as garbage during the sweep since the mark phase has completely traversed all active data structures. Such collectors cannot immediately be run so the mark phase and sweep phase execute concurrently.

According to one aspect of the present invention, there is provided a method of reclaiming memory space allocated to data structures comprising data objects linked by identifying pointers, comprising the repetitious cycle of the following steps: traversing the data structures to identify and mark active data objects; deleting data objects that were marked as garbage during the previous cycle; marking data objects not marked as active as garbage; and executing the finaliser of any data object marked as garbage having a finaliser.

The method and apparatus of the present invention permit marking and sweeping to be executed concurrently by explicitly marking garbage objects. By marking objects as garbage on one sweep phase and reclaiming those objects on the next, the deletion of objects is separated completely from the marking phase. Thus memory is reclaimed concurrently with marking live objects, improving the efficiency of garbage collection.

The step of executing the finaliser of any data object marked as garbage having a finaliser may comprise the step of queuing each finaliser to be executed and executing each queued finaliser in turn, and the step of traversing the data structures to identify and mark active data structures may include the step of traversing the data structure of each object having a finaliser queued to be executed and marking it and its descendants as active.

A first mark and a second mark may be used in alternate cycles to mark objects as garbage, wherein data objects having the alternate mark to that used in the current cycle are deleted.

A total of data objects having unexecuted finalisers may be compared to a total of data objects of active data structures having finalisers and, if the totals are equal all objects not marked as active are deleted.

The method and apparatus of the present invention efficiently reclaim memory allocated to garbage and finalisable objects as quickly as possible.

The data structure corresponding to a data object having a finaliser to be queued for execution may be processed to find any other data object of the data structure also having a finaliser queued for execution and, if the other data object is higher in any hierarchy of the data structure, the finaliser is

queued for execution after the finaliser for the other data object; otherwise, if the data object is higher in any hierarchy of the data structure than the other data object, the finaliser is queued for execution before the finaliser for the other object.

5 According to another aspect of the present invention, there is provided a data processing apparatus comprising a data processor coupled with a random access memory containing a data structure comprising data objects linked by identifying pointers, the processor being configured to provide the following in a repetitious cycle for operating on the stored plurality of data
10 objects: means for traversing the data structures to identify and mark active data objects, means for deleting data objects that were marked as garbage during the previous cycle, means for marking data objects not marked as active as garbage, and means for executing the finaliser of any data object marked as garbage having a finaliser.

15 The means for executing the finaliser of any data object marked as garbage having a finaliser may comprise means for queuing each finaliser to be executed and executing each queued finaliser in turn, and the means for traversing the data structures to identify and mark active data structures may include means for traversing the data structure of each object having a finaliser
20 queued to be executed and marking it and its descendants as active.

 The data processing apparatus may include means for obtaining a total of data objects having unexecuted finalisers and means for comparing the total with a total of data objects of active data structures having finalisers and, if the totals are equal deleting all objects not marked as active.

25 The data processing apparatus may include means for processing the data structure corresponding to a data object having a finaliser to be queued for execution to find any other data object of the data structure also having a finaliser queued for execution and, means for inserting the finaliser in the queue, wherein if the other data object is higher in any hierarchy of the data
30 structure, the finaliser is inserted after the finaliser for the other data object, otherwise, if the data object is higher in any hierarchy of the data structure

than the other data object, the finaliser is inserted before the finaliser for the other object.

5 The invention will now be described by way of example only, with reference to the accompanying drawings, in which:

Figure 1 is a block diagram of a data processing system suitable to embody the present invention;

10 Figure 2 represents the concurrency of marking and sweeping operations in prior art garbage collection systems;

Figure 3 is a flow chart of a memory reclamation method according to an embodiment of the present invention;

Figure 4 is a flow chart of a memory reclamation method according to another embodiment of the present invention; and

15 Figure 5 represents a finalisable data structure.

Figure 1 represents a data processing system, such as a personal computer, which acts as host for a number of software utilities which may, for example, configure the system as a browser for data defining a virtual environment. The system comprises a central processing unit (CPU) 10 coupled via an address and data bus 12 to random-access (RAM) and read-only (ROM) memories 14, 16. These memories may be comprised of one or several integrated circuit devices and may be augmented by a system hard-disk as well as means to read from additional (removable) memory devices, such as a CD-ROM. Also coupled to the CPU 10 via bus 12 are first and second user input devices 18, 20 which may suitably comprise a keyboard and a cursor control and selection device such as a mouse or trackball. Audio output from the system is via one or more speakers 22 driven by an audio processing stage 24. Video output from the system is presented on display screen 26 driven by display driver stage 28 under control of the CPU 10. A further source of data for the system is via on-line link to remote sites, for

example via the Internet, to which end the system is provided with a network interface 30 coupled to the bus 12.

The present invention is particularly embodied in memory management for a working area of the RAM 14 under control of the CPU 10; a controlling
5 program for this may initially be held in ROM 16 and loaded up with the operating system on power-up. The following examples are compliant with concurrent garbage collection and memory management techniques in the Java 7 (Sun Microsystems Inc) virtual machine and the terminology should be construed accordingly: it will be recognised however that the present invention
10 is not restricted to Java-compliant systems, nor to concurrent (or incremental) garbage collection, or to purely virtual memory management.

In the following, a number of particular terms having special meaning will be used, as follows. Root objects are globally accessible stored data objects in systems: in Java, these are represented by static data fields and
15 stacks. Reachable objects are stored data objects connected to the data structures accessible from the root objects of the system. During garbage collection, the marking process should find all reachable objects in one garbage collection cycle. Non-reachable objects are those stored data objects not connected to the data structures, they are inaccessible to running
20 programs and are hence candidates for garbage collection.

Finalisable objects are data objects which have nominally become garbage as they are non-reachable, but which also have a non-empty finaliser that has not yet been executed. A finaliser may only be executed once, so an object is considered finalisable if its finaliser has not yet been executed by the
25 garbage collector. Finaliser-reachable objects are a subset of the group of non-reachable objects which are reachable via pointers from finalisable objects: these objects should not be garbage collected as they may be resurrected or otherwise modified by the finaliser.

Mutator threads are threads of execution in an incremental garbage
30 collection system: the name comes from their effect of mutating the heap memory holding the stored data objects concurrently with the garbage collection threads, whilst a garbage collection thread is a dedicated system

thread whose purpose is to collect garbage objects and create more free space in the heap. A finaliser thread is a dedicated system thread which will run finalisation code for unreachable objects with finalisers.

5 The detection period is the time period between the last reference to an object being removed and its detection as an unreachable object. In mark-sweep collectors this is determined by the duration of the marking cycle. The reclamation period is the time between an object being detected as non-reachable and the point at which it is deleted. In mark-sweep garbage collectors, this is determined by the duration of the sweep cycle and the
10 presence of finalisable objects.

By way of illustration, it will be useful to first consider a commonly used technique for finalising garbage collection, used for example in the Sun Java Virtual Machine v1.0.2. The typical non-finalising garbage collection cycle can be described independently of implementation details as firstly the partition of
15 the heap into reachable and non-reachable sets of data objects, followed by garbage collection of the non-reachable set. Finalisation complicates garbage collection by adding a new stage to each garbage collection cycle, such as to comprise partition of the heap into reachable and non-reachable sets, as before; followed by partition of the non-reachable set into garbage and
20 finaliser-reachable sets; and finally garbage collection on the garbage set and finaliser execution on the finaliser-reachable set.

This lengthening of the cycle means that the lag between objects becoming non-reachable and their collection is made greater, with the consequence that a larger proportion of the heap is used up by garbage
25 objects and therefore a smaller area of the heap is available for new data items. Also, the strong dependencies between the stages in the cycles mean that most implementations will have poor concurrency characteristics, with large pauses in the mutator threads as the garbage collection algorithm proceeds. Considering that the number of finalisable objects in a typical
30 system is a small percentage of the set of active objects, the added overhead of finaliser support is disproportionately large. The state of each object in the system is identified by two variables M and F, where M may be Marked

(including any other mark states) or Unmarked, and F can indicate Finaliser or No Finaliser. An object having variable M = Marked is reachable and has been found by the marking process whilst M = Unmarked indicates that the object has not yet been found to be reachable.

5 The separate tasks relevant to the garbage collector and their interaction in the simple system are shown in Figure 2. The mutator threads MU may create new objects in the heap at any time. The marking stage MA involves the tracing of objects from root objects. Objects found by this tracing have M = Marked, and marking can be performed in the garbage collection
10 thread or, in the illustrated case of incremental garbage collection, in the mutator threads. The finaliser partitioning F-P involves sweeping over the heap, placing objects with M = Unmarked and F = Finaliser in a queue, then a mark phase using finalisable objects as roots: this runs in the garbage collection GC thread. The finalisation stage FS involves the execution of
15 finalisers in the queue, and sets F = No Finaliser for objects after execution. This may run in the garbage collection GC thread or, as shown, in a dedicated finaliser thread FS.

 This technique has some poor characteristics, the first of which is poor concurrency with the steps of mark phase MA, finaliser partition F-P, and
20 garbage collection sweep GC occurring serially. Consequently, the inherent independence between the tasks cannot be exploited in a multi-threaded system. Another problem is the generally long reclamation times for all objects. For non-finaliser-reachable garbage (which will be the vast majority), no garbage will be collected in this cycle until the finaliser partition F-P is
25 completed. This is wasteful as the objects will have been examined during the finalisation partition anyway. With incremental garbage collection algorithms, the system pause for mutators whilst finaliser partition F-P and the garbage collection GC sweep occur is unacceptably long, taking up two complete sweeps over the heap plus the finaliser-reachable marking phase.

30 Figure 3 is a flow chart of a memory reclamation method according to an embodiment of the present invention. A garbage mark variable (G) is associated with each object. Initially no value is attributed to the variable,

which is set when the respective object is identified as unreachable. A marking stage identifying and marking reachable objects is performed in step 300 using a technique such as those techniques discussed or referred to above. A sweep stage is then performed on the memory heap. Unmarked objects (i.e. those not reached in marking stage 300) are checked for an already existing garbage mark in step 310. If a garbage mark is found, the object is deleted in step 320. If no existing garbage mark is found, the object is assigned a garbage mark in step 330. If the finaliser variable associated with the object indicates it has an unexecuted finaliser, a reference to the object is added to the finaliser queue in step 340. The sweep stage comprises the assigning of garbage marks to unmarked objects, the deletion of data objects with existing garbage marks and the execution (or scheduling for execution) of finalisers. The marking stage includes the traversing of all objects reachable from root objects and marking them. For the purposes of the marking stage, any objects having an executing finaliser, or a finaliser scheduled for execution, is treated as a root object and therefore the object and its descendants are prevented from being reclaimed until the finaliser has finished its execution.

Steps 310-340 are repeated for each unmarked object in the memory heap. Once all of the heap has been processed, the finaliser for each object referenced in the finaliser queue is executed in step 350. This cycle thereby reclaims memory space allocated to those objects that had a garbage mark assigned during the previous cycle, and which remained as garbage even after execution (or scheduling of execution) of the finalisers during the previous cycle. The finaliser execution may have the effect of resurrecting the referenced object and/or any of its reachable descendants. A resurrected object would then be reachable from a root object. Once the finaliser has been executed, its reference is removed from the finaliser queue and the F variable of the object is set to F = No Finaliser.

During the next mark stage, any object having a garbage mark found to be reachable from a root object (as a result of finaliser execution) has its garbage mark cleared in addition to being marked. Preferably, finalising

objects, objects scheduled to be finalised and their descendants that are treated as root objects for the purpose of the mark stage do not have their garbage mark cleared so that they are reclaimed as soon as the execution of the finaliser has completed.

5 In the subsequent sweep stage all unmarked objects having a garbage mark, i.e. the garbage identified during the previous mark-sweep but not resurrected by a finaliser, can be immediately deleted, as explained above.

10 An embodiment will now be described as an enhancement to the method described above with reference to Figure 3, the embodiment providing a means for reliably detecting whether newly unmarked objects can safely be deleted (i.e. without waiting for the next mark-sweep cycle). This procedure is performed by tracking the number of objects with finalisers that have been found during marking, and comparing with a known and stored total allocated in the heap. As the finalisation queue is included in the marking process, the
15 only finalisable objects not located during marking are those that have become non-reachable during the previous collection cycle. When not all finalisable objects have been located in the previous mark phase, then there is a non-empty set of finaliser-reachable objects to identify in the sweep. However, when all finalisables are accounted for, all unmarked objects are known to be
20 garbage, and can be deleted immediately.

The object states are as for the Figure 3 example, but with three new global variables added:

 finaliser_count A count of the total number of finalisable objects in the heap

25 found_count The number of objects with finalisers located during the mark process

 shortcut A flag indicating how the sweep process should behave this cycle.

30 The tasks are carried out generally as for the Figure 3 example, but with some exceptions as illustrated by the flow chart Figure 4. In the case of the mutators, if a new object created has F = Finaliser then finaliser_count is incremented. During marking, if a found object has F = Finaliser, then

found_count is incremented in step 400. If, after all reachable objects have been marked, found_count = finaliser_count then no finalisable objects have become garbage this cycle and all unmarked objects can be safely deleted. The shortcut variable is set to True to signify this in step 410. Otherwise, the
5 shortcut variable is set to False to signify that at least there is one newly finalisable object to be processed this cycle and no newly unmarked objects should be deleted.

During sweeping, if the variable "shortcut" is true, then unmarked objects are deleted at step 430; if "shortcut" is false, then any unmarked
10 objects are processed as has previously been discussed with reference to Figure 3. Finalising is handled as before but, after running the finaliser thread for an object, F is set to No Finaliser, the object is marked, and both found_count and finaliser_count are decremented.

When the technique is applied to an incremental garbage collector as
15 described, the finaliser_count may change during the marking traversal. To make the counts consistent, the found_count must not include any objects found during the mark traversal or the comparison will yield the wrong result. In order to avoid this, during traversal found_count may have 1 subtracted if a marked finalisable object is encountered. Alternately, a second ongoing count
20 processed_count may be maintained to record the number of marked finalisers encountered during a sweep, with processed_count being deleted from found_count before comparison with finaliser_count.

Figure 4 shows this optimised finalisation model with additional reclamation through use of the "shortcut" variable: as will be recognised, in
25 comparison with the Figure 3 technique, that of Figure 4 guarantees 1-cycle reclamation when no finalisers are used and a maximum of 2-cycle reclamation for all non-finaliser-reachable data objects.

One variation to the methods presented above is to mark finalisers and their descendants as live during the sweep phase rather than waiting until the
30 next mark phase. In this way it is possible to tell exactly which objects are garbage at the end of the sweep phase rather than having to wait until the end

of the next mark phase. This would be of benefit if the mark and sweep phases were executed concurrently.

In the generally rare event that a data structure containing more than one finalisable object becomes non-reachable, the number of finalisables
5 detected this cycle becomes uncertain due to a race condition between the marking process and the progress of the sweep. For example, assume an object F1 is found by the sweep, added to the finalisation queue, and marked: F1 references another finalisable object F2, which is also non-reachable in this cycle. Depending on the behaviour of the marking algorithm, F2 may be
10 reached by the marker before or after the sweep in the next cycle reaches it. F2 is only added to the finalisation list this cycle if the sweeper reaches it first. A possible solution is to place objects on the finalising queue as they are marked, such that the same actions will be taken for an object regardless of whether it is found by the sweep or marking process.

15 An aggressive marking algorithm will tend to prevent lower finalisables in a connected data structure from entering the queue during the same cycle as the top levels of the structure, leading to a semi-topological ordering of finalisers. This typically aids recycling rates as higher objects in a data structure will have a larger reachable set than lower terms, and will therefore
20 tend to make more objects unreachable on deletion.

Conversely, marking of finalisables located during the sweep may be delayed until the sweep completes, such that all newly non-reachable finalisables will be added to the finalisation queue. This aids predictability, and permits a strengthening of the reclamation period definitions: 1-cycle
25 reclamation is guaranteed when no non-reachable finalisable objects are detected in the previous marking phase.

Figure 5 represents a finalisable data structure to be processed by a further enhancement of the method described with reference to Figures 3 and 4, in which the execution of finalisers of objects referenced in the finaliser
30 queue in step 350 is ordered.

The data structure 500 of Figure 5 comprises a number of data objects (510-570) linked by referencing pointers. Data objects 510, 550 and 560 have

finaliser methods associated with them. At a point during the execution of a program process, the data structure is no longer needed and becomes unreachable. Due to the fact that the memory heap holding the data objects is not ordered, the order in which the sweep finds the data objects is unpredictable. Supposing that data object 560 is found before 510, the order of the finaliser queue will be 560, 510 and 550. This order is not desirable as data object 560 is reachable from the finaliser of data object 510 and could therefore be resurrected during the execution of the finaliser for data object 510, negating the need to execute its own finaliser.

Therefore, once all objects have been processed by the sweep, the finaliser queue is ordered to provide a topological execution of finalisers. A breadth-first search is performed by following the pointers from each object referenced in the finaliser queue to its descendants. If a data object referenced earlier in the finaliser queue is found to be reachable from another data object referenced in the queue, the queue is re-ordered so that data object will have its finaliser executed before its descendant.

In this manner, the queue of finalisers for Figure 5 would be re-ordered to 510, 560 and 550.

The ordering of finalisable data objects referenced in the finaliser queue could alternatively be done as references to data objects are added to the queue, searching the queue to determine where to insert the reference.

In the embodiments described above, a binary state variable (G) is used to identify garbage. However, a multistate variable could preferably be used. For example, states A and B may be used, both signifying objects which are garbage. On one sweep phase, garbage is marked with State A, but is not deleted. At the end of the sweep phase any finalisers of finalisable objects are executed. On the next sweep phase, State B is used to mark garbage. On marking an object with State B, if it already has a State A mark it can safely be deleted as it is still garbage and has not been resurrected by a finaliser.

Although defined principally in terms of a software browser implementation, the skilled reader will be well aware than many of the above -

described functional features could equally well be implemented in hardware or a combination of hardware and software.

From reading the present disclosure, other modifications will be apparent to persons skilled in the art. Such modifications may involve other
5 features which are already known in the design, manufacture and use of data processing and/or data network access apparatus and devices and component parts thereof and which may be used instead of or in addition to features already described herein.

CLAIMS

1. A method of reclaiming memory space allocated to data structures comprising data objects linked by identifying pointers, comprising
5 the repetitious cycle of the following steps: traversing the data structures to identify and mark active data objects; deleting data objects that were marked as garbage during the previous cycle; marking data objects not marked as active as garbage; and executing the finaliser of any data object marked as garbage having a finaliser.

10 2. A method as claimed in claim 1, wherein the step of executing the finaliser of any data object marked as garbage having a finaliser comprises the step of queuing each finaliser to be executed and executing each queued finaliser in turn, and wherein the step of traversing the data structures to
15 identify and mark active data structures includes the step of traversing the data structure of each object having a finaliser queued to be executed and marking it and its descendants as active.

20 3. A method as claimed in claim 1 or 2, wherein a first mark and a second mark are used in alternate cycles to mark objects as garbage, wherein data objects having the alternate mark to that used in the current cycle are deleted.

25 4. A method as claimed in any one of the preceding claims, wherein a total of data objects having unexecuted finalisers is compared to a total of data objects of active data structures having finalisers and, if the totals are equal all objects not marked as active are deleted.

30 5. A method as claimed in claim 2, wherein the data structure corresponding to a data object having a finaliser to be queued for execution is processed to find any other data object of the data structure also having a finaliser queued for execution and, if the other data object is higher in any

hierarchy of the data structure, the finaliser is queued for execution after the finaliser for the other data object, otherwise, if the data object is higher in any hierarchy of the data structure than the other data object, the finaliser is queued for execution before the finaliser for the other object.

5

6. A data processing apparatus comprising a data processor coupled with a random access memory containing a data structure comprising data objects linked by identifying pointers, the processor being configured to provide the following in a repetitious cycle for operating on the stored plurality of data objects:

10

means for traversing the data structures to identify and mark active data objects;

means for deleting data objects that were marked as garbage during the previous cycle;

15

means for marking data objects not marked as active as garbage; and

means for executing the finaliser of any data object marked as garbage having a finaliser.

7. A data processing apparatus as claimed in claim 6, wherein the means for executing the finaliser of any data object marked as garbage having a finaliser comprises means for queuing each finaliser to be executed and executing each queued finaliser in turn, and wherein the means for traversing the data structures to identify and mark active data structures includes means for traversing the data structure of each object having a finaliser queued to be executed and marking it and its descendants as active.

20

25

8. A data processing apparatus as claimed in claim 6 or 7, wherein a first mark and a second mark are used in alternate cycles to mark objects as garbage, wherein data objects having the alternate mark to that used in the current cycle are deleted.

30

9. A data processing apparatus as claimed in any one of the claims 6 to 8, including means for obtaining a total of data objects having unexecuted finalisers and means for comparing the total with a total of data objects of active data structures having finalisers and, if the totals are equal deleting all
5 objects not marked as active.

10. A data processing apparatus as claimed in claim 7, including means for processing the data structure corresponding to a data object having a finaliser to be queued for execution to find any other data object of the data
10 structure also having a finaliser queued for execution and, means for inserting the finaliser in the queue, wherein if the other data object is higher in any hierarchy of the data structure, the finaliser is inserted after the finaliser for the other data object, otherwise, if the data object is higher in any hierarchy of the data structure than the other data object, the finaliser is inserted before the
15 finaliser for the other object.

11. A program storage device readable by a machine and encoding one or more programs of instructions for executing the method steps of a specified one of claims 1 through 5.

1/4




FIG. 1




FIG. 2

2/4




FIG. 3

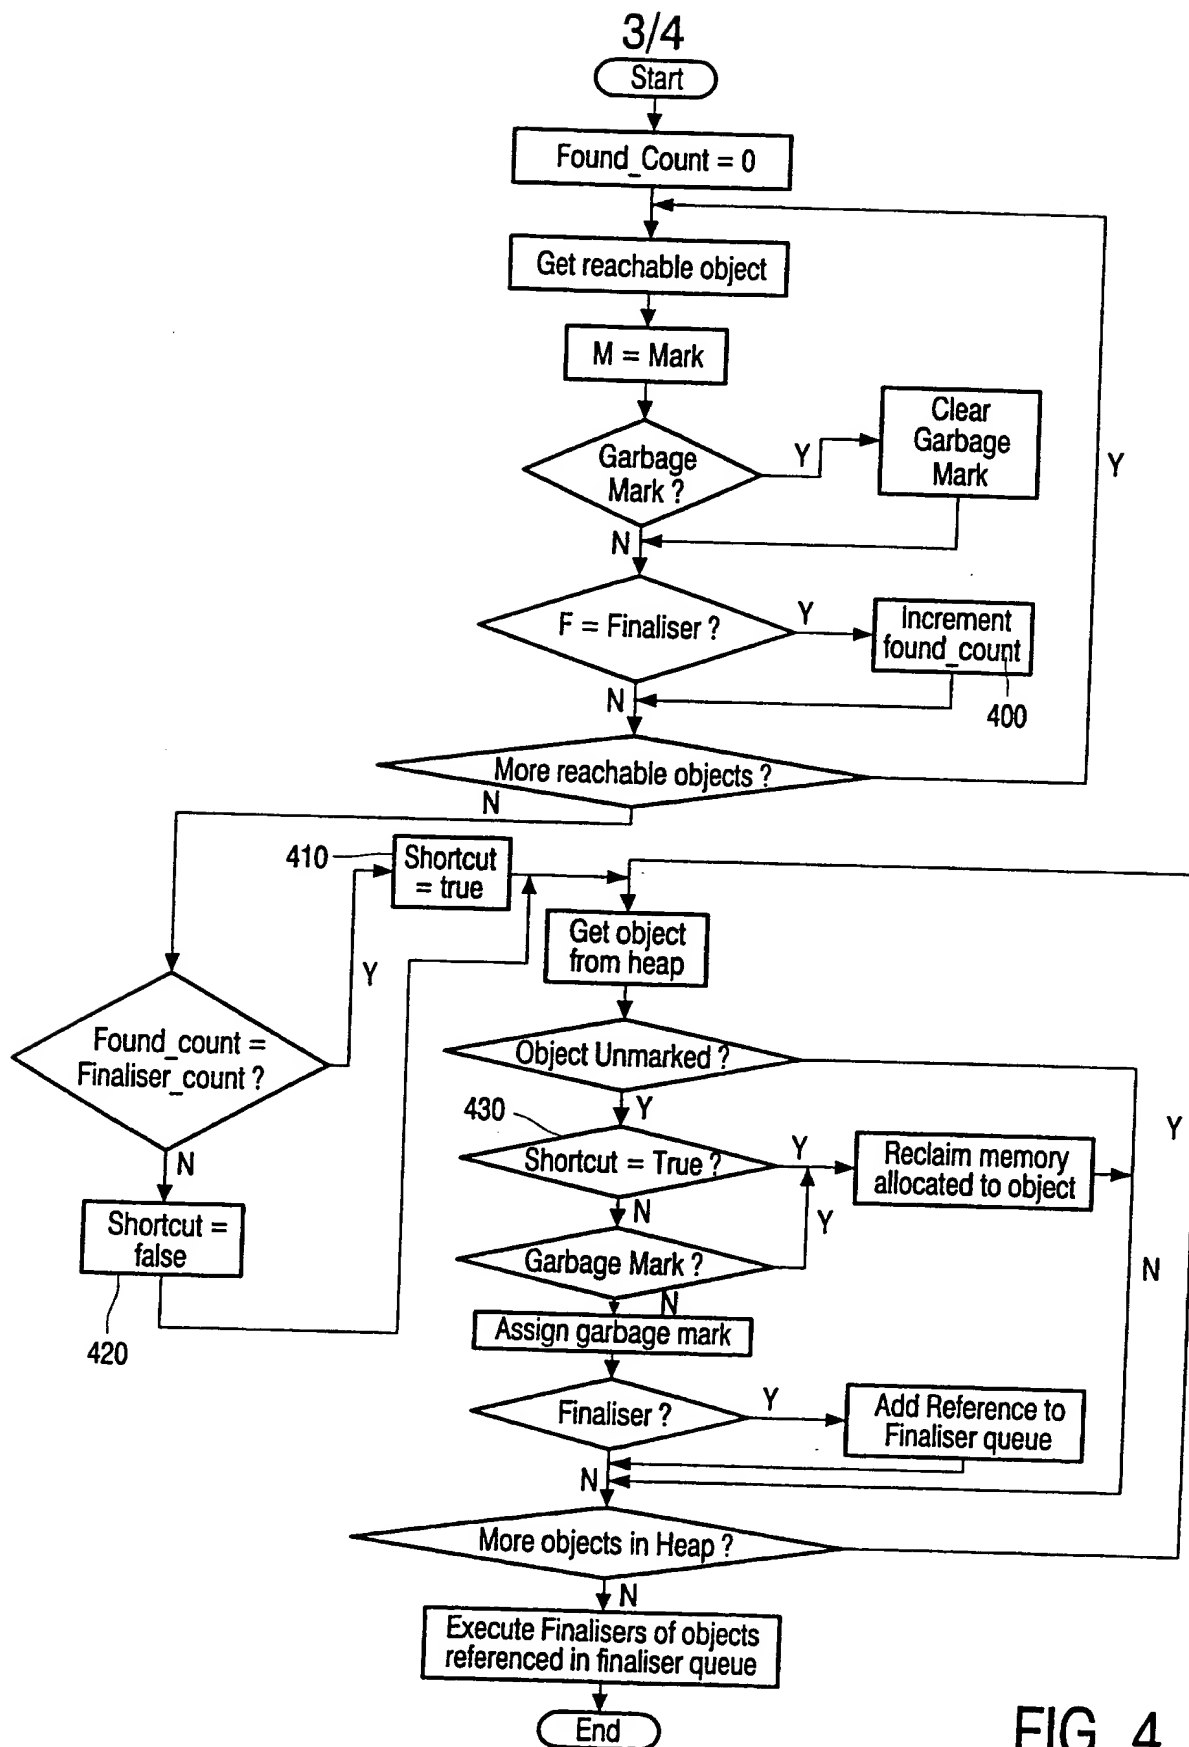


FIG. 4

4/4




FIG. 5

INTERNATIONAL SEARCH REPORT

Intern. Application No
PCT/EP 00/02094

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F12/02

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	ATKINS M C ET AL: "THE ACTIVE DEALLOCATION OF OBJECTS IN OBJECT-ORIENTED SYSTEMS" SOFTWARE PRACTICE & EXPERIENCE, GB, JOHN WILEY & SONS LTD. CHICHESTER, vol. 18, no. 11, 1 November 1988 (1988-11-01), pages 1073-1089, XP000110495 ISSN: 0038-0644 page 1078, line 20 -page 1086, line 14	1-3, 5-7, 10, 11
A	US 5 274 804 A (JACKSON FRANKLIN R ET AL) 28 December 1993 (1993-12-28) column 4, line 30 -column 6, line 46	1, 2, 5, 6
A	US 5 392 432 A (VANDENDORPE JAMES E ET AL) 21 February 1995 (1995-02-21) column 25, line 10 -column 27, line 9	1, 2, 6, 7

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

13 July 2000

Date of mailing of the international search report

21/07/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Nielsen, O

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/EP 00/02094

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5274804	A	28-12-1993	CA 2123994 A	27-05-1993
			EP 0642686 A	15-03-1995
			JP 7507887 T	31-08-1995
			WO 9310504 A	27-05-1993
US 5392432	A	21-02-1995	NONE	

NO. 18 384 304

THIS PAGE BLANK (USPTO)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☒ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (USPTO)